

DAVIS REVIEW

NGT 00200

MAY 25 1988 PM



Ted Nelson
Director
Xanadu
8480 Fredricksburg, Suite 138
San Antonio, TX 78229

Date: May 25, 1988

To: Ted Nelson
Xanadu/Autodesk


From: L. Mills Davis
Chairman, Expert Communication'88

Subject: Conference overview

I have written a 20 page overview of hypermedia and artificial intelligence in electronic publishing which I am enclosing for your review. The purpose of the report is to introduce the conference and to provide attendees something in writing that they can walk away with. Each attendee will receive a copy when they register.

I look forward to getting together with you at Expert Communication'88! To reach me before the conference, please call me at 202-667-6400. I plan to fly to San Jose on June 5th.

Sincerely,



L. Mills Davis

DAVIS REVIEW
Published by DAVIS INC

Business Office:
2704 Ontario Rd NW
Washington, DC 20009

Telephone: 202-667-6400
Telex: 5106018770DAVIS

Publisher: L Mills Davis
Manager: Brigitte Reyes
Design: Ann Chaparos

Production Systems:
Page layout—Aldus PageMaker 3.0.

DAVIS REVIEW is published 10 times a year at 2704 Ontario Road NW, Washington, DC 20009 by DAVIS INC. All foreign and US rights reserved. ISSN 0893-3510. Title registered in US patent Office. Yearly subscriptions are \$300 in US, \$325 in Canada and Mexico, \$350 other than North America.

1988 Copyright DAVIS INC.
All rights reserved.

Important: This report contains results of proprietary research. Reproduction in whole or in part is prohibited without written permission of DAVIS REVIEW.

Executive summary

A new wave of publishing technology is beginning to build. It comes from the convergence of three technologies: hypermedia, artificial intelligence (AI), and electronic publishing. This new hybrid we are calling "hyperpublishing" technology.

This issue of DAVIS REVIEW examines the significance of hyperpublishing technology as an evolutionary direction in electronic publishing. We discuss the nature of pioneering activities in various industry segments and identify some of the practical uses of hypermedia and artificial intelligence in electronic publishing today. We explore how far we can go with today's hyperpublishing technologies, giving examples of vendors and products, and discuss likely future developments.

The impact of hyperpublishing technology over the next decade will be broadly felt in four ways. It will (1) redefine the nature of tools for authoring, design, and communication arts, (2) transform publishing processes, (3) introduce new media and new forms of publishing, and (4) engender new (global) literary environments for research, collaborative authoring, publishing, and use of information.

But is it real? Practical applications of hypermedia and artificial intelligence in electronic publishing are an established fact in the marketplace. The scope of these applications is impressive: speech recognition, text scanning, graphics recognition, natural language interfaces, hypermedia and expert systems for authoring and design, idea processors, intelligent search and retrieval, precis writers, automatic document generation, hypermedia publications, expert systems for communication, expert systems for printing and publishing management, networked literary environments, and much more.

Hyperpublishing technology, however, still has a long way to go. In many respects, it is very much a pioneer's market. Applications have succeeded that could live within definite technology constraints.

The outlook is for rapid development of hyperpublishing technology. Market interest is growing, especially now that PCs make viable platforms for AI- and hypermedia-based applications. The number of hyperpublishing-related projects, and the level of investment in them by corporations, government agencies, and commercial publishing organizations is rising sharply. The number of vendors developing hypermedia- and AI-based tools for publishing and the pace of new product introductions is also increasing.

It is time we learned more about hyperpublishing technology and what it means for our industry.

L. Mills Davis

Hyperpublishing

Hypermedia and artificial intelligence in electronic publishing

Contents

Hyperpublishing technology 1

Hypertext 2

Artificial Intelligence 3

Expert systems 3

Supercomputing 4

Neural networks 4

Impact of hyperpublishing technology 5

Industry activity 6

Technology directions 7

Content acquisition 7

Speech recognition 7

Text scanning 8

Handwriting recognition 9

AI-based image scanning 9

AI-based line art scanning 9

Graphic recognition 10

Natural language interfaces 11

Verballization 11

Hypermedia authoring 11

Idea processing 12

Document parsing 12

Computer-assisted editing 13

Computer-assisted translation 13

Computer-generated extracts and precis 13

Computer advisory dialogues 14

Automatic content generation 14

Visualization 15

Design automation 15

Hypermedia art and design tools 16

Visual expert systems 16

AI-based image enhancement 16

Computer-assisted charts and diagrams 16

AI-based typography 17

Publication knowledgebase management 17

Desktop knowledgebases 17

Automatic indexing 17

Electronic librarians 18

Intelligent search and retrieval 18

Visual knowledgebases 18

Hyperdocument generation 18

Smart templates 18

Intelligent document formatting 19

Intelligent text retrieval 19

Hypermedia publications 19

Expert communication systems 19

Speech synthesis 20

Publication/printing management 20

This white paper examines hyperpublishing technology—the practical use of hypermedia and artificial intelligence in electronic publishing.

The questions it seeks to answer are:

What is hyperpublishing technology?

Why is it important?

What impact will it have on publishing?

What industry activity is ongoing?

What are the practical applications?

What future developments are likely?

Hyperpublishing technology

The convergence of hypermedia, artificial intelligence, and electronic publishing technologies is producing an important new direction in publishing that we call "hyperpublishing." We define hyperpublishing as the use of hypertext/hypermedia, smart computers, and networks to create, produce, and distribute information in a visual (or audio-visual) form.

Hyperpublishing synthesizes three concepts: *non-linear media* (hypermedia) and *embedded intelligence* (artificial intelligence), and *preparing documents at (or near) typeset quality by computer* (electronic publishing).

What is hypermedia? Ted Nelson, one of the pioneers of hypertext and hypermedia, defines hypermedia as non-sequential (or non-linear) text, graphics, and imagery. Hypermedia is a hybrid that combines a database method, a knowledge representation scheme, and a human interface modality. The concept is quite simple: windows on a display screen are associated with objects in a database, and links are provided between these objects.

What is artificial intelligence? Intelligence is the power to acquire and apply knowledge. Artificial intelligence (AI) is intelligence embedded in machines.

AI attempts to make electronic publishing systems more useful by making them handle publishing tasks in a way that, if people were solving these problems, you would call these people intelligent. Four key areas of AI have importance to electronic publishing:

Speech—recognition and synthesis of human speech

Vision—recognition and rendering of text, graphics, and imagery

Natural language—understanding and generation of spoken, written and visual communication

Expert systems—knowledge representation, learning, reasoning, planning

What is electronic publishing? Electronic publishing is a generic process that uses computer and electronic imaging technologies to acquire, develop, produce, and disseminate information in a visual form. Electronically produced publications typically contain text, graphics, and continuous-tone images. Newer publishing media, such as CD-ROM and CD-I, can also include audio, motion video, and animation in addition to static visual content.

Hypermedia, AI, and electronic publishing are approaches or tools for publishing applications. They are not applications by themselves.

Hypertext

According to Jeff Conklin (MCC, Austin, TX), hypertext implementations can be grouped into four broad categories:

Literary systems
Problem exploration tools
Browsing systems
General hypertext technology

Literary systems are large on-line libraries in which interdocument links are machine supported. These macro literary environments are intended to support research, publishing, reading, collaborative work.

Problem exploration tools support unstructured thinking on a problem when many disconnected ideas come to mind. Applications include authoring, outlining, problem solving, and design.

Browsing systems are like smaller scale literary systems where ease of use is crucial. Applications include teaching, reference, public information dissemination.

General hypertext technology includes systems for reading, writing, collaboration, etc.

A hypertext system is a database, semantic network, and a human interface modality.

The basic features are as follows:

The database is a network of textual and graphical nodes or hyperdocument. You can directly access data, which differs from traditional use of queries.

Windows on the display screen correspond to nodes in the database on a one-to-one basis. Each has a name or title which is shown on the window. Only a small number of nodes are ever "open" (as windows) on the screen at the same time.

The human interface supports standard window system operations: for example, you can reposition, resize, close

windows or put them aside as small window icons. Closing a window causes it to disappear after any changes that have been made are saved to the database node. Clicking with the mouse on the icon of a closed window causes the window to open instantly. The position and size of a window or icon (perhaps also its color and shape) are clues to remembering the contents of the window.

Windows can contain any number of link icons which represent pointers to nodes in the database. A link icon usually contains a short textual field to suggest the content of the node it points to. Clicking on a link icon with a mouse causes the system to find the node referenced and to open a new window for it on the screen.

The user can create new nodes and new links to nodes easily. This facilitates annotating, commenting, elaborating on ideas, etc. as well as establishing new connections to existing nodes.

The user can browse the database by (1) following links and opening windows successively to examine their contents, (2) searching the network (or part of it) for some string, keyword, or attribute value, and (3) navigating around the hyperdocument using a browser that displays the database structure graphically. The user can select whether the nodes and links display their labels or not.

Hypertext authoring and reading environments have these advantages:

Ease of tracing references
Ease of creating new references
Both hierarchical and non-hierarchical information structuring
Global views
Mixing global and local views
Customized documents
Modularity of information
Consistency of information
Multiple paths of inquiry
Ease of collaboration

Hypertext systems also have disadvantages.

Getting "lost in hyperspace" can be a problem. You can lose your sense of location and direction in a nonlinear document. With conventional text you can search forward or backward to find material. Hypertext offers more degrees of freedom (more dimensions in which you can move). It has more potential for becoming disoriented.

Graphical browsers, information filters, and query/search mechanisms are the main ways of coping with disorientation. Browsers represent the hyperdocument in a 2D or 3D spatial environment providing a kind of information map. Information filters elide information so that the user is presented with a manageable level of complexity and detail, and can shift the view or suppress details while navigating through the database. Query techniques apply boolean operations and some combination of keyword search, full-string search, and logical predicates.

Cognitive overhead is another problem. Additional effort and concentration is necessary to maintain several tasks or trails of association at one time. For example, a new idea occurs to an author while writing. He presses a button. An empty window appears on screen. He records the new thought. He presses a button and returns to the previous point he was writing. The overhead comes with having to deal explicitly with the mechanisms for linking ideas: for example,

What is the name of the new idea in the database?

What is a one-line summary for it to assist browsing?

What type of link does the new idea have to the previous one?

What type of links does the new idea have with other topics?

Artificial intelligence

Intelligence is the power to acquire and apply knowledge. Artificial intelligence is intelligence embedded in machines.

The term "artificial intelligence" was coined in 1956 when ten scientists convened a conference at Dartmouth College to discuss the development and applications of artificial intelligence. They predicted, among other things, that by 1982 "computers do all the work and we will be playing." They assumed erroneously that intelligent behavior could be readily copied onto computer programs.

Marvin Minsky in Society of Mind argues that intelligence is not a fixed commodity, but a process.

The idea is that intelligence can arise from the interconnection and interaction of non-intelligent components or "agents." These agents have no intelligence of their own but perform simple functions when activated by other agents. Collections of agents may function together as an "agency." Higher-order agencies will exhibit increasing amounts of intelligence, intelligence being the ebb and flow of activity in a network of agencies, sub-agencies, and agents.

Minsky's theory serves as a strategy for useful AI-based systems. Simple, clearly defined algorithms and logical functions (already implemented in computer systems) fit nicely with Minsky's concept of agents. As parallel processing (and neural computing) advance, more difficult problems of connectivity between agents can be addressed. What remains, however, is to develop a deeper understanding of human thought processes.

Expert systems

According to Roger Thibault (Artificial Intelligence Today) "An expert system is a computer program that incorporates the knowledge of an expert (or group of experts) on a particular subject and manipulates the data we input in a fashion that 'mimics' the human reasoning process." The expert system contains logical rules and steps similar to those that an expert uses to solve a particular problem. What sets expert systems apart from other forms of programming is the ability to do inexact reasoning from heuristic data or rules-of-thumb.

Expert systems have three primary elements:

- (1) a knowledge base or rules
- (2) an inference engine which is used to order the rules, manipulate data, and apply the rules to the current problem (or set of facts)
- (3) the problem, which usually comes in the form of data supplied at the time the expert system is run on the computer.

Expert systems are not something radically new, just a different way of defining and implementing a solution. According to Herbert Schorr of IBM "Knowledge engineering (building expert systems) allows you to handle new, tough problems that are too costly or too painful to do with conventional programming techniques."

Commercially applicable expert systems appeared in the 1980s. Eighteen months ago, however, scarcely a handful of these systems existed in business and government. Today, an estimated 1000-3000 expert systems are in daily use. Industry sources say that the number is increasing by 50 percent annually.

Expert systems are entering a period of transition.

Once companies were satisfied with experimental, stand-alone expert

systems built with special hardware and software. Now they want to make expert systems accessible to large numbers of people and integrated into conventional technologies. As the industry realigns, companies that relied on special-purpose machines are languishing.

PC-based AI is becoming an important segment of the market.

William Martorelli (New Science Associates, Norwalk, CN) estimates that in 1987 about \$70 million was spent for PCs to develop and deliver expert systems. He notes that Compaq Computer Corporation (Houston, TX) estimated that about 3 percent of its Deskpro 386 PC computer sales in 1987 were pointed toward expert system activity, or around 3000 units. Carl Wolf (Gold Hill Computers, Cambridge, MA) adds "There's not much you can do on a Lisp machine that you can't also do on a 386."

According to Dan Shafer, associate editor for PC AI, the Mac-based expert system tool market stands today about where the PC market did 18 months ago. A choice of tools exists today. Significant commercial expert system activity is likely within the next 12 to 18 months.

Apple's HyperCard provides a very easy to use language for developing hypermedia applications called HyperTalk. Expert system tools are being offered now that link to the HyperCard/HyperTalk environment.

Cognition Technology's (Cambridge MA) MacSMARTS, for example, interfaces to graphic, spreadsheet, and database files as well as to HyperCard. You can launch MacSMARTS from inside HyperCard and you can link any rule to HyperCard and have it open a specific stack. As of this report, no automatic exchange of data between HyperCard and MacSMARTS is supported.

Supercomputing

According to W. Daniel Hillis (designer of the Connection Machine) "Supercomputing is an enabling technology for artificial intelligence. Just as you couldn't build an airplane without first developing engines powerful enough to drive them, you can't build artificial intelligence without fast supercomputers."

Supercomputing pursues speed and raw processing power.

Sales of supercomputers, which cost \$5-25 million each, have averaged around 25 percent annual growth over the past decade. The installed base is currently around 300 machines. Supercomputer sales are expected to top \$1 billion in 1988.

The leaders in the supercomputer market are Cray Research (Minneapolis, MN) and Control Data Corporation (St. Paul, MN). Japanese manufacturers (NEC, Hitachi, Fujitsu) represent about one-fourth of the market. The Japanese now build the fastest single-processor supercomputers.

IBM is giving major backing to a parallel-processing supercomputing venture headed by Steve Chen (formerly of Cray Research). The goal of this effort is a machine 100 times as fast as any currently on the market. AT&T's Bell Laboratories is also introducing a new parallel-processing computer this year.

Much of the supercomputing research is funded by the U.S. government.

Most supercomputing research is directed toward defense and intelligence uses. The Department of Defense spent several hundred million dollars last year to step up the speed of the fastest machines. Presidential Science Advisor, William Graham has recommended that Congress appropriate an additional \$1.7 billion to support development of parallel-processing supercomputers that by the mid-1990s could crunch data at teraFLOPS speed.

Neural networks

Neural networks are a new form of parallel computing with applications in speech recognition, text and handwriting recognition, and graphics/image recognition.

The idea behind neural networks comes from theories developed to explain how the brain, neurons, and synapses accomplish cognitive tasks so rapidly, with individual elements that operate so slowly.

Since most cognitive processes take no longer than a few hundred milliseconds while neurons operate in milliseconds (not nanoseconds as do conventional computers), the brain must perform its processing feat in no more than 100 serial steps through massive parallelism.

According to James House (San Diego, CA), the most popular neural-net design is the Rumelhart Net, named after designer David Rumelhart, a researcher at the University of California in San Diego. The Rumelhart net sandwiches one or more "hidden slabs" between an input slab and an output slab in order to geometrically multiply the signal pathways between the two.

Massive parallelism affords neural networks a high degree of fault tolerance (ability to recover from processor failure), associative recall (ability to retrieve information very rapidly based on content), and graceful degradation (the ability to guess if there is no exact match for the requested information).

How do neural networks learn to recognize letters, graphics, or to synthesize speech?

To teach the neural net to recognize a new letter, for example the letter [A], the user writes a cursive letter [A] on the digitizing pad a few times. As the letter is written, its form triggers some neurons and not others in the input slab. Each triggered neuron sends a signal down each of its pathway connections to the inner (hidden) layer

of neurons. Inner layer neurons receive signals from some, but not all input layer neurons. If the sum of the signals received is sufficiently great, the inner neuron sends a signal down to the next layer of neurons. The propagation of signals eventually forms a pattern of stimulated output neurons that matches the preprogrammed pattern of the letter [A].

The neural net uses a process called "back propagation" to train itself to recognize the letter.

It does not successfully recognize a letter the first time. The neural-net retraces the signal routes back from the output layer to the input layer, adjusting the signal level that will trigger neurons to fire and adjusting the strength of the signals that these neurons send down the pathways to the next layer. The back propagation process is repeated until the adjustments (or weights) cause the input signals to generate the preprogrammed letter pattern on the output layer. After a few more samples of the handwritten letter (and additional weightings have been determined) the neural net will recognize almost every instance of the letter [A].

The NETtalk voice-synthesis program developed at Princeton University dramatically illustrates this learning process. NETtalk changes letters into simple English words and sentences. You can hear the net adjusting weights as it trains itself to speak. It starts out talking gibberish (like an infant), then learns to speak right before your eyes.

Are there any practical neural network products for electronic publishing applications?

Not many. This technology is too new, probably 2-3 years away from serious commercialization. We estimate that no more than 10-15 companies are working on neural network-based products today. Some are working on applications. Most are working on hardware.

Impact of hyperpublishing technology

Hyperpublishing will affect the tools, methods, and end-products of publishing.

The impact will be broadly felt:

First, as it redefines electronic publishing tools

Second, as it transforms the publishing process

Third, as it introduces new media and new forms of publishing

Fourth, as it engenders new global environments for research, authoring, publishing, and use of information.

Hyperpublishing technology can provide a much richer and more integrated creative environment for authors, illustrators, and designers than is provided by current word processors, drawing programs, and page make-up systems.

Embedding intelligence in devices, workstations, and networks promises to reshape the publishing process as we know it, compressing steps, providing new capabilities, improving quality, and creating new publishing environments that leverage the productive efforts of authors and audiences alike.

Hypermedia combined with expert systems technology expands our concept of what a publication is. It is a new kind of printing press that produces personalized, interactive (non-sequential) publications with embedded simulations, animations, and explanations. For example, expert systems are being used in industry to train, provide quick-reference to published material, assist in maintenance, and provide feedback for revision and updating product literature. The result is a new form of technical publishing that offers improved economics for documentation, training, and support throughout a product's life cycle.

Hyperpublishing technology applied globally will produce literary environments with new forms of information access, distribution, and usage. Already, several universities and government agencies are undertaking pilot programs to engineer electronic libraries and integrated research/publishing environments where much of the information that a scientific researcher needs and uses will be readily available and easily retrieved, independent of its physical location.

Why is hyperpublishing technology becoming important now?

Hyperpublishing technology is a megatrend whose time has come. Its roots go back at least thirty years. It indicates a massive evolutionary step forward, rather than a revolutionary break with the past.

Pioneering work has been done over the past three decades in each of the component technologies. Today's developers are building from this base of applications experience.

A threshold in price-performance has been reached enabling more rapid development. Hardware will no longer be the limiting factor that it has been. The fact that AI and hypermedia approaches can be implemented on standard platforms (with, if needed, off-the-shelf add-ons to provide specialized processing capabilities) means that they will evolve as mass technologies and become ubiquitous. Within two years, next generation workstations and personal computers will have ten times the power of today's machines. As Bill Poduska of Stellar Computer (Newton, MA) remarked, "By 2001 we'll have machines with 100 MIPS of computing power and give them to our kids for Christmas."

Hyperpublishing technology meshes with the strategic concerns of major corporations. Businesses and govern-

ment agencies are seeking competitive advantages, new capabilities, and operational efficiencies that the new approaches can confer. They are funding projects.

Hyperpublishing technology also aligns with the strategic interests of major computer manufacturers. The announced direction, products, and research commitments of companies such as Apple Computer, IBM, AT&T, and consortiums such as Microelectronics and Computer Technology Corporation (MCC) underscores the importance of hypermedia and artificial intelligence as core technologies for future products. According to John Sculley of Apple Computer,

"Today we are on the verge of creating new tools which, like the (printing) press, will empower individuals, unlock worlds of knowledge, and forge a new community of ideas. The tools that show the most promise for the new learning environment build on these core technologies: hypermedia, simulation, and artificial intelligence." (MacWorld Expo 1988, speech)

Finally, hyperpublishing technology is in accord with the strategic interests of the nation as a member of the world community. Programs such as Japan's Fifth Generation project make this abundantly clear. As Time Magazine stated in a recent article (2/28/88), "The country that leads the world in supercomputing and artificial intelligence will hold the keys to economic and technological development in the 1990s and beyond."

Industry activity

Who is interested in hyperpublishing technology? why? and what pioneering activities are under way?

The total potential market for electronic publishing solutions is huge. Almost every organization publishes, processes documents, trains, and manages information in a visual form. Out of a universe of 5 million corporations employing 25 million white collar workers, we estimate that the addressable market for electronic publishing systems is probably around 1.5 million sites. We estimate that businesses will invest \$50-75 billion in electronic publishing solutions over the next decade.

Although hyperpublishing technology will pervade electronic publishing by the end of the decade, it is not clear that a distinct market for "hyperpublishing solutions" will develop.

If pressed for an estimate, we think it is reasonable to project that hyperpublishing technology (meaning hypermedia and AI embedded in electronic publishing solutions) will develop into a billion dollar market in the 1990s.

How real is it today? For years the idea and the promise of hypermedia and artificial intelligence has greatly exceeded its actual performance. We began our research for this report concerned that hypermedia might be all hype and that AI might not be able to do anything of value for publishing. We found, instead, that they are becoming commercial technologies with practical publishing applications and real products.

Pioneering applications of hyperpublishing technologies are underway in many different kinds of organizations.

We estimate that somewhere in the range of \$500 million currently is being invested in R&D, and prototype applications.

Universities are devoting resources to developing prototype electronic libraries, scholar workstations and tools for translation, scientific visualization, and hypermedia based tutoring.

The vision which motivates this work is of a macro literary environment in which information, independent of physical location, can be readily available, and easily retrieved. As Dana Scott of Carnegie Melon University (Pittsburgh, PA) commented, "After decades of high expectations, we are on the brink of a time when much of the information that a scientific researcher needs and uses can be made available on a workstation."

Information industry companies have been investigating hyperpublishing technologies for use in leveraging investments in content databases and information delivery services.

These companies are exploring speech recognition and speech synthesis, advanced search and retrieval services, computer advisory dialogues to assist users, automatic abstracting, precis writing, and document generation, hypermedia publications, personalized publications, new media, and expert systems-based information delivery.

The Department of Defense (DoD) and advanced manufacturing organizations are exploring hyperpublishing technologies as part of a move into an all digital environment for technical information.

Hyperpublishing technologies will provide the base for engineering, publishing, training, maintenance, and information dissemination activities.

For example, a voice-activated, notebook-size interactive maintenance system could provide (1) ongoing training using hypermedia-based audiovisual tutorials, (2) ready-reference through a

hypertext reader, demand printing, or speech synthesis, (3) real-time assistance with field diagnostic and test procedures using simulation-driven expert systems, and (4) feedback for problem resolution and updates through voice or handwritten annotations transmitted to product maintenance. In DoD, around \$500 million has been allocated for research and development of prototype "new information products." Current plans call for full deployments in the early 1990s.

Financial services have been innovators in AI-based information services.

Current projects include expert systems for portfolio management and loan processing as well as natural language programs to extract and summarize text, and to generate forms and documents such as personalized insurance policies.

Commercial printing and publishing businesses traditionally have been early adopters of new publishing technology.

The development of editorial front-ends in newspapers, for example, lead to the explosion of word processing in the office. Today firms working with the commercial publishing industry are working on speech recognition, intelligent scanning of pictures, fast search and retrieval of large scale databases, new media publications, and expert systems to layout pages and configure presses.

Government and corporations are investing in varied projects involving hyperpublishing technologies.

These activities include expert system to explain policies and regulations, hypermedia publications, "friendly" front-ends to databases, and AI-based systems to generate documents such as forms, policies, proposals, specifications, and contracts.

Technology directions

What publishing problems can hyperpublishing technologies address and how far can we go with today's products?

In the following discussion, it is important to remember that we are dealing with emerging technologies, not mature ones. The technology directions, products, and applications that we discuss in this report usually deal with just pieces of a publishing solution, not the whole thing. The successful products we found do something specific and useful for applications that can live within definite constraints. This caveat should be kept in mind.

AI and hypermedia have a broad range of electronic publishing applications.

For purposes of this report, we group these applications into the following categories:

- Content acquisition
- Verbalization
- Visualization
- Knowledgebase management
- Hyper-document generation
- Publishing/printing management

Content acquisition

Content acquisition is concerned with inputting text, graphics, and imagery from spoken, written, printed, photographic, and electronic sources.

Most input to electronic publishing systems comes from a keyboard. The use of digitizers and scanners is growing. The direction we are heading is toward tools that can capture text, graphics, and image information directly from original sources, provide comprehensive interfaces to external information sources, and support massive input of textual information into the personal and corporate knowledge bases of the 1990s.

AI-based approaches are being used to:

- Recognize speech/convert it into text
- Recognize typewritten or typeset text
- Recognize symbolic expressions
- Recognize tabular material
- Understand handwriting or proof marks
- Enhance the quality of imagery
- Recognize graphic content in pictures
- Assist users to provide information

Speech recognition

Speech recognition converts sound into text. Commercial speech recognition applications in electronic publishing exist today, but they have a ways to go before they reach their full potential.

As computers come to understand speech, an author, editor, artist, or designer might not have to use the keyboard or digitizer to enter commands or input text. The direction we are heading is toward systems that can automatically transcribe natural language conversation. Large vocabulary speech recognition systems for personal computers should reach the market by 1990.

Speech recognition capabilities are classified in two ways.

First, these systems are either speaker-dependent or speaker-independent. Second, the systems are capable of recognizing either isolated-words, connected-speech, or continuous-speech. The trend is toward speaker-independent, continuous-speech recognition.

Speaker-dependent systems require the user to train the system to recognize their voices by making templates of their voice patterns and enrolling them in the system. Speaker-independent word recognition is phonetically based and uses AI-based feature analysis rather than template matching to identify words.

Isolated-word systems identify individual words. Pronounced pauses must separate each spoken word or vocabulary item. Connected-word systems recognize sets of words uttered without pauses; however, pronounced pauses must separate utterances. Typical capacity is about 30 seconds. Continuous speech recognition requires no pauses but may entail limited vocabulary.

Reaching the goal of speaker-independent, continuous-speech recognition requires on the order of 10 times more computational power than connected-word systems.

The challenge of speech perception is that error and ambiguity permeate all levels of the speech decoding process. Algorithms must distinguish speech from noise, apply knowledge of both language and context to the speech recognition task, and accommodate gracefully the missing phonemes, noise, diverse environmental conditions, variability due to speaker, and incomplete or inaccurate knowledge-source representations. Although much closer today, we are several years away from unlimited vocabulary speech recognition systems.

Practical speech recognition systems are available today.

For example:

Speech Systems (Tarzana, CA) offers a 20,000 word, speaker-independent, continuous speech recognition on a general-purpose platform. Their approach uses phonetic recognition based on pattern-recognition, linguistics, signal processing, and AI. Spoken words are translated into phoneme codes and the phonetic codes are translated into words on the workstation (or host computer).

Kurzweil AI (Waltham, MA) markets a 1,000 word, speaker-dependent isolated speech recognition system called Voicesystem. Still under development is a 5,000-10,000 word, voice-activated word processor for dictation called Voicewriter that uses linguistic expert systems. For example, Nash General Hospital (Rocky Mountain, NC) uses VoiceEM, a Kurzweil speech recognition computer, in the emergency room to assist physicians and nurses to fill in forms.

IBM (Yorktown Heights, NY) has developed (and shown in television advertisements) an experimental 5000 word, speaker-dependent, isolated word, voice-to-text transcription machine that uses stochastic modeling techniques to determine (based on probability) what word is meant when words sound similar to one another. IBM discovered, however, that a 5,000 word vocabulary is simply not enough for the office and is increasing the vocabulary to 20,000 words.

Dragon Systems (Newton, MA) offers a 1000 word, speaker-dependent, continuous-speech recognition system that uses stochastic processing techniques.

Votan (Fremont, CA) makes small vocabulary (150 word) speaker-dependent and speaker-independent continuous-speech recognition systems with keen accuracy and high noise tolerance. They also package a voice terminal that provides voice output and voice recognition for applications such as voice mail and direct voice input (e.g. quality control inspections).

NEC (Melville, NY) markets a 250 word, speaker-independent, isolated-word speech recognition and playback systems as well as a 150 word, speaker-dependent, continuous-speech recognition system.

Text scanning

Text scanning of already typed or typeset documents uses optical character recognition (OCR) techniques to produce text files in machine readable format.

OCR technology has had commercial applications in electronic publishing for several years. The industry leaders include Datacopy (Mountain View, CA), Dest (San Jose, CA), Kurzweil (Waltham, MA), and Palantir (Santa Clara, CA). Each firm provides products with strong optical character recognition (OCR) capabilities and good levels of accuracy.

Desktop text scanning technology is moving toward recognition of complex document formats containing varied typefaces and sizes, symbolic expressions, tables, graphics, illustrations.

The direction we are headed is toward fully automatic capture and conversion of the resulting database into a useful knowledge representation (document structure, hypertext links, indexes) or a tagged document structure. Another

direction for the 1990s is high-speed scanners that can convert massive amounts of text-based information into commercial, corporate, and academic knowledgebases.

Current OCR technology has limitations.

The range of type faces and sizes varies with the price of the scanner, but is generally narrow. Small typefaces (such as 6 point type) do not scan with acceptable error rates. The scan rates, while faster than a typist can keyboard, are far too slow for high volume production scanning.

Intelligent text scanning technology has important future applications in electronic publishing as a tool for building free-text knowledge bases.

A great deal of information of critical importance during the research and authoring phase of publishing currently falls outside the scope of office automation for lack of cost-effective ways to enter and maintain it in an on-line database. Such information includes magazine, journal, and newspaper articles, excerpts from books and published reports, business correspondence and so on. As organizations expand their use of optical (and other high density) storage systems to capture this information, intelligent text scanning will become a key technology for building comprehensive archives of textual knowledge.

AI-based text scanning approaches are currently being used to recognize composite documents with symbolic expressions, tabular material, and graphics as well as text. Avalanche's (Boulder, CO) Intelligent Markup System, for example, uses intelligent scanning and machine-learning techniques to automatically mark-up documents. The system detects

relatively simple structures (headings, paragraphs, bulleted lists, numbered lists, row/column tables, rules, and images) and inserts corresponding tags (or macros) into the output file.

Handwriting recognition

Intelligent scanning techniques are being used to understand handwriting, written comments, and proof marks. Applications exist for script recognition in electronic publishing, but there are limits to current technologies.

Computer understanding of handwriting or cursive script would be beneficial in several publishing contexts. For example:

An author could input a handwritten manuscript without keyboarding, either by using a scanner, or by writing it out on a digitizing tablet.

An editor could "red-line" and annotate a manuscript or proof by hand (thus providing a clear audit trail), scan the marked up pages into the computer, and have the computer apply these handwritten editorial corrections to a wordprocessor file.

A user in an east Asian culture could have a simpler way to interact with computers. In Japan and China, for example, the large number of graphic symbols that make up the writing system makes authoring, editing and publishing by computer very difficult to do with a keyboard.

The direction that we are heading is toward new types of terminals that combine features of digitizing tablets and video displays.

Readers may remember that astronauts in the film "2001" recorded data by hand in notebook-size computers.

The data entry field and display screen were one and the same, so that numbers, graphics, or text written on the notebooks could be automatically read, interpreted, and recorded.

Such a notebook may be realized in the next 12-18 months according to Dave Fenwick, president of Airus Inc. (Portland, OR). His company is developing a handwriting recognition package with this type of device in mind.

Another way of thinking about this technology is as an intelligent drawing table for designers, draftsmen, and technical illustrators. At least one Japanese manufacturer, Seiko Instruments, has developed a prototype LCD tablet and stylus which it calls the ViewPad.

AI-based image scanning

AI-based image scanning is being used to optimize the quality of printed color images.

In previous issues of DAVIS REVIEW (Vol 1. No. 1 New Directions in Color, and Vol. 1 No. 2 Emerging Standards for Electronic Color Publishing and Prepress), we examined some of the constraints of the four-color printing process. Scanning of continuous-tone images for four-color printing demands compression of the tonal range apparent in the original transparency to what can be achieved with inks and paper. There are many ways to compress the image and all of them affect the quality of the color reproduction. A skilled operator is able to adjust the scanner controls and tone compression curves to achieve the best color reproduction possible given the features of the original image and the characteristics of the inks, paper, and printing press being used.

AI-based image scanning uses expert systems technology to emulate the decisions made by a skilled scanner operator.

Scitex (Bedford, MA) introduced the first "smart scanner" for use in high-end color printing applications. Its expert system has knowledge about tone compression, color correction, color separation, halftoning, under color removal, gray component replacement, inks, presses, papers, and other aspects of the color printing process. It also accepts feedback to calibrate and fine tune the color separations that it produces.

AI-based scanning technology will become a major direction in the desktop scanner market over the next 12-18 months.

Applications include color separation, image enhancement, features extraction, and image compression.

AI-based line art scanning

Line art is a form of illustration prepared with high-contrast (usually black and white) graphic elements. Unlike continuous-tone imagery, line art can be reproduced by conventional printing without halftone screening.

AI-based techniques are being used to convert images into line art (line conversion) and to store scanned line art in a compressed form.

Imagin (Karlsruhe, West Germany), for example, has developed a program for storing line art in a compressed form. It scans an image line by line, and records only the brightness/color changes. Once converted, technical drawings, text, and artwork are resolution-independent. They can be scaled and printed at very high resolution retaining the sharp edges and high contrast of the original image.

Graphic recognition

Graphic recognition converts an image into a graphic figure by identifying and linking the collections of points, lines, arcs, and parametric formulas that describe the geometry of an object.

Graphic recognition technologies have been applied most to technical publishing and computer-aided design and drafting (CADD).

Users of CADD systems need to be able to capture and access databases of existing drawings in a form that they can revise and incorporate into new designs or publications.

Source drawings are complex and vary in legibility and precision. The drawings may be unscaled or scaled, with no dimensions, few dimensions, or all dimensions. Sizes range from very large (E-size drawings) to very small (aperture cards) with large format drawings very common.

Raster-to-vector conversion is sometimes used when merging a hardcopy illustration into a technical publication database or loading a drawing into a CADD system.

The graphics recognition process is normally divided into four parts: pre-processing, vectorization, recognition, and post-processing.

Pre-processing provides some level of drawing restoration before proceeding to vectorization. This entails simple manipulation of image data to remove dirt, smudges, discolorations, or to fill in broken and faded lines. The program examines the scanned image, detects line segments and attempts to connect them together.

Vectorization is the process of converting the bit-map dots or pixels provided by the scanner into vectors or line representations.

Recognition is the ability to detect common shapes such as circles, arcs, boxes, or a variety of curved shapes. Once a particular shape has been recognized, its characteristics can be recorded and stored very efficiently. For example, a circle will be recorded simply as a center and a radius value which requires much less storage than recording all of the dots that make up a raster circle.

Post-processing, or clean-up, corrects for inaccuracies inherent in the raster-to-vector conversion. For example, the points where lines, arcs, and circles meet are cleaned up. Similarly, line weights can be adjusted. The program endeavors to recognize graphic symbols, text, hand annotations, and line art content of existing drawings in a way that both preserves the connectivity between different graphic, text, and image elements and maintains accurate dimensionality. The resulting graphics file format must be compatible with existing CADD software.

Graphics recognition is cost-effective for some applications, but, significant limitations still exist.

For example, graphics recognition tends to produce a very large number of vectors. Even with software to condense the number of line segments, considerable manual editing is usually required. Throughput is measured in terms of hours per drawing, not drawings per hour.

Some companies offering graphic recognition systems for engineering drawings include:

Autodesk Inc. (Sausalito, CA) offers a program which converts technical drawings (reflective art) scanned on a desktop scanner into 2D geometric line drawings. This process attempts to recognize points, lines, arcs, and symbols in the image and to reconstruct the figure as a display list.

ANA Tech (Littleton, CO) markets a scanning and vectorizing system of A-to E-size drawing input into CAD and technical publishing systems.

Formtek (Pittsburgh, PA) makes a raster-scan and editing system for direct input of raster drawing images into CAD systems.

Industrial Vision Systems (Lowell, MA) produces a large-format document scanner and subsystem for vectorizing drawings.

Metagraphics (Woburn, MA) makes an AI-based scan/edit system for drawing input to a CAD database.

Optigraphics (San Diego, CA) offers a complete scan/edit workstation for input to CAD systems.

Scan-Graphics (Broomall, PA) manufactures scanners, conversion/editing software, and automatic (CAD) input systems.

Scitex (Bedford, MA) sells a scan/edit system for drawing conversion and input into CAD, publishing, and mapping applications.

Skantek (Warren, NJ) offers an automatic digitizing scanner system

Sysscan (Dearborn MI) markets a complete scan/edit system.

Tektronix (Beaverton, OR) produces a complete scan/edit system for drawing input to CAD applications.

Graphics recognition in other areas of publishing is a somewhat different application.

Graphics recognition for office and commercial publishing demands a high degree of aesthetic control. Illustrations for commercial publications tend to have a richer style and graphic

vocabulary than is found in engineering drawings or (most) technical illustrations where the use of typography, line weights, textures, tonal range, and color is much more generic and restricted.

Users of electronic art, design, and publishing systems need to be able to capture and access databases of existing artwork in a form that they can revise and incorporate into new illustrations and publications. This need is shared by commercial firms that manufacture clip art collections, graphic templates, symbol collections, and other types of picture libraries intended for use by electronic publishing systems.

The direction we are headed is toward sophisticated, low cost, raster-to-vector conversion of graphic arts illustrations.

A few companies are working on graphics recognition software for graphic arts illustrations. For example, T-Maker (Mountain View, CA) is developing software to scan and convert clip art and illustrations into PostScript files.

Desktop scanner manufacturers are likely to address this application incrementally over the next few years.

Natural language interfaces

How can we make it easier for people to access a computer system and get at the information they want?

Natural language interfaces allow someone who does not know the syntax of a specific operating system or query language to talk with the computer to obtain information. The interface attempts to interpret the request on its own, solicit additional information from the user to help it interpret parts that it did not under-

stand, rephrase the request in terms that the database system will understand, and request verification that it has correctly interpreted the research request.

Cognitive Systems (New Haven, CN), for example, develops computer advisory dialogues for public information dissemination, point-of-sale information, database access, and expert systems applications. These systems have built-in a knowledge base for the subject-domain to enable them to interpret questions put to them by users.

Natural language interfaces are being used to assist users to provide information and fill out forms.

Mind Path Technologies (Dallas, TX), for example, offers a product called Enform that enables designers to build "smartforms" consisting of input screens, instructions, and explanations for complex instructions, and on-line cross field editing. Applicants or persons preparing the form can receive various levels of help at different times while filling out the form.

The direction we are heading is toward intelligent information gatherers.

The system reads a form, whether paper or electronic. Its questions are posed verbally to the user. Verbal replies are converted to electronic answers and relayed to the data bank.

Verbalization

Verbalization is concerned with supporting authors and editors in the development of written material and with generating certain types of document content automatically.

Hypermedia- and AI-based approaches are being used to:

- Author in a hypermedia environment
- Process ideas and outlines
- Parse text and marking-up a document
- Check spelling, grammar, and syntax
- Check word use, style, and readability
- Translate text into another language
- Generate precis, summaries,
- Extract text
- Generate explanations
- Generate document content

Hypermedia authoring

Today, authoring tools are, for the most part, wordprocessors and form letter generators while editing tools are little more than spelling checkers. We need tools that support research, idea processing, and design aspects of writing, manage quantities of text-based information and provide more of the expertise of a human editor in software to assist an author. The direction we are heading is toward hypermedia authoring environments with tools that exploit windowing environments, multitasking operating systems, and networked file servers.

Hypermedia offers new ways for authors and designers to work as well as new possibilities for accessing large or complex information sources.

If authoring is viewed as a word- and sentence-level activity, a word processor is an adequate tool. However, few authors simply pour out finished text. Authoring involves structuring of ideas,

ordering presentation, and exploring concepts. Editing involves more than word smithing and polishing.

Authoring is the design of document content.

The proper tool for this level of authoring is an idea or concept processor. Hypertext authoring tools can express ideas as nodes in the database and windows on the screen. The writer thinks of new ideas, opens a window and develops the thought. Later the writer can link them to existing ideas, or leave them isolated if it is too early to make such associations. The hypermedia environment assists the refinement of the unstructured network of ideas into a final polished document.

The convergence of word processing, desktop publishing, and hypertext is inevitable.

The appearance of general hypertext technology in the form of Apple's (Cupertino, CA) HyperCard, Xerox's (Palo Alto, CA) NoteCards, Brown University's (Providence, RI) Intermedia, Tektronix's (Beaverton, OR) Neptune, MCC's (Austin, TX) PlaneText, and Owl's (Bellevue, WA) Guide gives some indication of breadth of hypermedia development that has already occurred. The interest that companies such as Microsoft (Bellview, WA), Lotus (Cambridge, MA), and Aldus (Seattle, WA) are showing in hypermedia based authoring environments may signal something of the direction that the industry will take in the future.

What is a hypermedia authoring environment like?

Jay Bolter (UNC, Chapel Hill, NC) and several colleagues have developed a hypertext-based authoring system for the Macintosh called Storyspace that permits document reading and writing in several dimensions at once. It is much more than an outline processor.

Writers can organize thoughts in many more ways, throwing boxes up on the screen, typing text into them, and then drawing lines to show associations. It is more like brainstorming. Outline processors force the writer to organize thoughts in a progression. Storyspace supports hierarchical structure in a more flexible way. Creating "paths" enables you to link nonsequential thoughts. These features also permit Storyspace to be used as a reader for non-sequential fiction.

Idea processing

One of the strengths of the hypermedia environment is its ability to process ideas, outlines, and hyperdocuments rather than just words.

There are current products on the market that offer problem exploration tools to support unstructured thinking and concept exploration. For example:

MaxThink (MaxThink Inc., Piedmont, CA) is an outline/idea processor that can handle 4096 topics in 99 levels. Topics may be text or graphics. MaxThink provides commands for traversing the outline vertically and horizontally as well as commands for reorganizing, adding to, deleting from, copying, or cutting-and-pasting content. The prioritize, binsort, randomize, levelize, categorize, tag/untag, and fence commands provide other ways to move topics around and are useful for conceptual thinking and evaluating ideas.

Houdini (MaxThink Inc., Piedmont, CA) is a tool for making linkages between topics. It can handle 2,500 topics and 7,500 links, provides fourteen ways of linking text or graphics, eight ways to display link relationships, and eight ways to print them out. For example, Houdini can automatically link topics by text patterns, keywords, associations, comparisons, divisions, or boolean relationships.

A number of PC-based outline processors let you move back and forth from writing, outlining, revising without changing programs. For example:

Living Videotext's Thinktank
Ashton-Tate's Framework
Applied Microsystems' daVinci Idea Processor
Travelling Software's Idea!
Intelligent Environment's Deja Vu
Seaside Software's askSam
SoftWorks Development's PC-Outline

Document parsing

AI-based authoring tools exist that can parse text and mark-up a document automatically.

In electronic publishing, text tagging schemes are used to describe the logical structure of written material and to convey specific information about how it should be composed. Tagging makes it possible to differentiate document semantics (such as headings or paragraphs) from specific typesetting instructions (such as typeface and size), so that a document can be composed in more than one way without recoding the document. SGML (Standard Generalized Markup Language) is the most important markup language for developing tagged text structures.

As interest has grown in text markup and composition based on style sheets, so has interest in products to assist the markup process.

SoftQuad (Toronto, Ontario, Canada) has developed an authoring tool that automatically assigns predefined SGML tags to text based on its understanding of the text being input.

Sobemap (Brussels, Belgium) is developing a suite of AI-based document handling tools which it calls DOC-IT. DOC-IT has five parts:

TELL-IT is an expert system used to create document type definitions as defined in SGML.

MARK-IT is a full featured SGML parser that reads text files according to the document type definition and inserts tags.

ED-IT is a SGML smart editor that allows documents to be created and edited according to a specific document type definition.

SET-IT is an AI-based formatter for SGML documents that uses roughly 20 heuristics to guide the formatting process along with various constraint relaxation techniques.

PRINT-IT is an intelligent printer driver.

Avalanche's (Boulder, CO) Intelligent Markup System inserts SGML tags into scanned text based on a document profile provided by the user.

Computer-assisted editing

Computer-based editing tools being used to check spelling, syntax, grammar, and writing style.

The goal of AI-based editing tools is to perform the task of an experienced editor, detecting errors in spelling and grammar, and suggesting ways to rephrase passages of text to make them more understandable and to make them conform to the patterns of high quality language usage.

In practice, other than spelling checking, the contributions of current "style and usage" editing tools are much more modest.

AT&T Bell Laboratories' Writers Workbench was the first of the computer-aided editing tools. In addition to spelling checking, the Writer's Workbench provides feedback about word usage and syntax and applies readability formulas.

Language Systems Inc. (Providence, RI), has developed several tools for computer-aided editing, including a phonetically-based spelling corrector, an electronic thesaurus, spelling warning systems (capitol vs capital), and a syntax-based grammar correction system. Henry Kucera, LSI's founder, is professor of linguistics and cognitive sciences at Brown University.

Houghton-Mifflin's CorrecText is a sentence-based grammar checker. Because it analyzes the grammatical construction of sentences, it is able to pick up many of the confusable words (such as who/whom, its/it's etc.) that show up commonly in documents. It has a built-in spelling checker. In addition, it can perform seven different readability tests (including Flesch, Spache, and Dale-Chall).

Computer-assisted translation

Machine translation converts documents written in one natural language to corresponding documents written in another language, but with equivalent meaning.

Early attempts in language translation were essentially sophisticated dictionaries and failed miserably. One famous example was "The spirit is willing, but the flesh is weak." Translated into Russian and back into English the statement became "The wine is good, but the meat is spoiled."

Creating fluent translations by computer remains an elusive goal. However, progress has reached the point where practical AI-based translation products and services exist.

For example:

Smart Communications (New York, NY) has developed an AI-based translation system that is being used by the Canadian Immigration Authority

to translate documents from English into French and from French into English.

Automated Language Processing Systems (Provo UT) offers systems for computer-aided translation to and from several European languages.

High Technology Solutions (Poughkeepsie, NY) has a system for translating English into Chinese.

Weidner Communications (Northbrook, IL) makes multilingual wordprocessors with bi-directional translation capabilities.

Computer-generated extracts and precis

AI-based products are being used to extract textual information and to generate precis, abstracts, and summaries.

Products that can extract relevant textual information from simple documents such as natural language telexes are available.

For example:

Sobemap's (Brussels, Belgium) TEXTTRACT extracts textual information from financial transfer notices between banks, or police fine notices, etc. The program scans for keywords in the text. It uses them to understand the context, extract information, and insert tags into the text.

Cognitive Systems' (New Haven, CN) has also developed a product that extracts essential banking and commercial transaction data from telexes.

Carnegie Group (Pittsburgh, PA) offers a product called Language Craft that extracts textual information from simple documents through parsing and case framing techniques.

Some products exist that can prepare abstracts, summaries, or précis of free-text information.

For example:

Brattle Research Corporation (Cambridge, MA) offers both a product and a service for searching and recompiling information. This system reads text and limited types of charts (e.g. Wall Street Journal), extracts textual information, builds a database, and writes a report to a specified format. The Seattle Police Department uses this system to prepare summaries of news stories.

MIT's NewsPeek project carried this idea a step further. The user would tell the agent (the system) about her interests and preferences. The system would search news wires and on-line databases for these topics, obtain stories, prepare extracts and summaries, perhaps adding a few things on its own, and deliver to the user a "personal newspaper".

Personalization of information is a trend in publishing that will grow.

Commercial information retrieval services based on the MIT agent concept are beginning to appear. The customer prepares a profile that identifies his/her interests and the database service monitors databases and other sources to prepare and deliver a customized newsletter.

Computer advisory dialogues

Computer advisory dialogues can be used to help define document types, audience profiles, and graphical styles.

Xpercom's (Carrollton, TX) Thoughtline, for example, is like talking to a friend about the speech, report, or presentation that you are preparing. It uses pattern matching techniques to ask you a series of questions that help you to organize your thoughts. For instance, it may ask, "What is the main idea?" or "What incident or analogy can you use to set the stage?"

The questions change as the program helps you brainstorm. It asks follow-up questions to help you clarify main points, and probes for examples, incidents, and documentation that provide background material. (You can also add your own questions.) After you've answered all of the questions, Thoughtline gives you an outline which lists your audience, goals, opening, thesis, key points and supporting data, and ending.

Automatic content generation

AI-based approaches are being used to generate document content from knowledge bases.

Automatic generation of document content is a major direction in corporate, technical, and commercial database publishing.

Automatic content generation seeks to translate information from the formal language in a computer's memory into ordinary language. The resulting document might be a single sentence (answering a query) or a technical manual with extensive text with illustrations.

Different approaches to automatic generation of document content exist, but share some common characteristics.

They all seek to leverage intellectual content in some way. They tend to deal with a well defined business problem in a fairly narrow subject area and the communications required are also well defined. The applications that we are seeing today include forms, policies, contracts, specifications, proposals, and technical literature.

Corbel & Co. (Jacksonville, FL), for example, provides two document generation services that are driven by expert systems. Both of these applications produce time and dollar savings.

Corbel's first service involves preparing pension plan submissions to IRS.

Lawyers go through an interactive dialogue with the system which defines the key parameters of the plan. The document generation process draws on a knowledgebase of legal language provided by experts in pension plans, editing and adapting it to the particular situation. The benefit for lawyers using the system is that they can accomplish two weeks work in about two hours. The document produced is considered a great first draft. The final product can be viewed interactively and edited before being composed and typeset for submission to the IRS.

Corbel's second service involves preparing architectural bid specifications. Most of the effort in an architectural design is spent on getting the drawings completed. The specification document, which is the legal document, is often left to the last minute. An interactive dialogue captures the key information for the document. The knowledgebase contains language pertaining to bid specifications. The document generation process understands what to include, what to leave out, and how to adapt and modify the legal language based on the content provided by the customer. After editing, the final specification is composed and typeset by the system.

Automatic document content generation techniques can also be applied to documents intended for the general public.

In this case, the design of the document and the quality of the language database become critical. Siegel and Gale (New York, NY), for example, developed a database and graphic design approach for the Chubb Group to support preparation and printing of customized insurance policies. The goal was to provide customers with a simple, clearly written, easily understood policy document and to do away with the reams of boiler plate and legalese that makes most insurance policies unintelligible. The database contained all of the definitions, clauses, and terms rewritten into plain English. The graphic design solution was

DAVIS REVIEW

developed to facilitate automatic document content generation in a mainframe environment with printing on a Xerox 9700.

Expert systems are being combined with CAD systems to speed bid/proposal preparation.

The savings on engineering man hours, particularly for mature products whose basic design is generally accepted, can be impressive. For example,

Wisdom Systems (Chagrin Falls, OH) developed a system for Babcock & Wilcox to automate preparing bids for large centrifugal fans. The process used to take several days of an engineer's time because each one needed to be customized to its application. Today, the fan design criteria are in the system, so that they just input the specifications and in ten minutes get back a whole proposal package.

Corporations are expressing interest in automatic generation of technical literature from product databases.

AndWare Publishing's (Milford, NH) VIDA is an AI-based product for generating technical literature running on Symbolics (Cambridge, MA) equipment. The method they describe is simple. First, VIDA parses existing content files such as product specifications, engineering memoranda, CAD drawings, etc. into units of knowledge (which it calls knots). Information describing the form of a publication (its layout and typography) and its context (document profile, audience, subject, scope, etc.) are similarly parsed and added to the knowledgebase. Upon demand, VIDA extracts and recombines this information and transforms it into reader-friendly product literature. As of this writing, we have not visited a VIDA installation and cannot comment on the quality of the results they achieve.

Visualization

Visualization is concerned with supporting authors, illustrators, and designers in the development of visual material and with generating certain types of visual content automatically.

Visualization is a method of computing that transforms the symbolic into the geometric. It embraces both image understanding and image synthesis.

Hypermedia- and AI-based approaches are being used to:

Automate design

Design and illustrate in hypermedia

Correct and enhance graphics and images

Generate charts, diagrams, models, maps, and synthetic imagery from text files, spreadsheets, and databases

Generate typefaces

AI-based imaging and vision has been a subject of interest since the 1960s. Pattern recognition, for example, has been used to help diagnose diseases by comparing X-rays and CAT scans against healthy models. It has also been used in military surveillance to help distinguish potential targets in a landscape. Actually, there is only a slight use of graphics in most of these applications. In the future, though, we will see AI-based imaging applications that both remember what they "see" (the features that they extract) and draw the picture.

Design automation

Although links between artificial intelligence and computer graphics are currently limited, AI-based graphics will emerge as an important technology for design-automation applications.

According to Carl Machover (Machover Associates, White Plains, NY), AI-based graphics software currently represents only about \$10 million of the \$2 billion graphics software market. Expert systems, however, are already beginning to be used in design-automation. For example:

Computervision's (Bedford, MA) Personal Architect automatically creates double-wall plans and 3D models from a single line schematic and technology file.

ICAD's (Cambridge, MA) mechanical engineering system incorporates a knowledge-base of design rules that help designers avoid potential design errors and structural interferences.

Graphic Horizons' (Cambridge, MA) CADD expert system checks architectural designs against national fire codes.

Hypermedia art and design tools

The direction we are heading is toward hypermedia environments for art and design that support visual idea generation, features extraction, picture editing, and visual database management.

Current visualization tools for electronic publishing consist of separate drawing, charting, painting, and page make-up software packages. The new environments will include expert systems for visualization as well as design. A key product focus will be to embed the expertise of illustrators and designers into "smart" tools to assist authors to visualize information content such as charts and diagrams or other illustrations acceptably with a minimum of effort.

Visual tools are beginning to appear that support design and illustration in a hypermedia environment.

Meta Software Corporation's (Cambridge, MA) Design, for example, is a graphical outliner that can be used to create flow charts, simulation and analysis diagrams, systems designs, information networks, computer programs, technical documentation, project planning, production line processes, organizational charts, and presentation graphics.

A diagram consists of graphical objects (such as circles, ellipses, boxes, images, etc.), text associated with each object, and connectors that relate the objects in various ways. Objects stay connected when moved. Hierarchical (parent-child) linkages between objects are implemented through hypertext-like display conventions. You highlight the object and select the external reference (child) to display the hidden content (for example, a detailed explanation, or a detailed schematic). To return, you select the *parent* option.

Visual expert systems

The direction that we are headed is toward expert systems for visual information management.

The visual expert system would scan photographs, video, and graphics and extract visual features. It would express these attributes in a textual, graphic (or symbolic) form and convert them into a useful knowledge representation or tagged image structure.

Bernise Glenn (Glenn Associates, Santa Monica, CA) points out that visual classification is the central problem. We need to define the aesthetics, or significant visual attributes for the application as well as the language to use in describing them.

Next, comes the problem of graphics recognition or telling the AI program how to go about recognizing visual features when it sees them. Finally, comes the problem of features modification or how to revise visual attributes when the users says "make the eye-brows bushier" or "make that wall less textured" or "raise the roof 3 centimeters." Designers and illustrators using visual expert systems will probably work from thumbnail sketches just as they do now. The power of the new tools will become apparent as they edit and explore variations.

AI-based image enhancement

AI-based techniques are being used to correct and enhance graphics and images.

Recognition Technology, Inc. (Holliston, MA), for example, develops hardware for image processing and machine vision applications using AI machine vision techniques. Applications include motion detection, object detection through smoke, fog, clouds, etc. using intelligent filtering.

An image that looks pretty is not necessarily an image that will reproduce well in print or other media.

The direction we are heading is toward visual expert systems for photo retouching, color editing, and rendering that will become part of continuous-tone illustration, graphic design, and page makeup software packages.

Computer-assisted charts and diagrams

AI-based approaches are currently being used to generate charts, diagrams, models, maps, and synthetic imagery from text files, spreadsheets, and databases.

In business and scientific visualization, the trend is to package expertise to help user construct the right chart, map, or diagram for a particular purpose and to express it with appropriate aesthetics (graphical style).

One area that is receiving a lot of attention right now is "presentation intelligence" or expert systems for desktop presentations.

Some computer-assisted diagramming packages use AI-based features to simplify and speed construction of complex information graphics.

The user makes a rough sketch of the diagram by hand. To digitize the drawing, the user clicks over a spot on the diagram to assign a symbol from a predefined template. When the symbols have been placed, the user goes back and connects symbol pairs by pointing, clicking, and indicating the type of connection or other attribute. When all the connections have been made, the user signals the computer to layout the drawing. The computer uses rule-based design and constraint-relaxation heuristics to set type, place symbols, set line styles, routes, dimension the drawing, etc.

Default chart types and graphic styles are techniques used in some statistical graphics packages.

SPSS's (Chicago, IL) Graphics, for example, is a statistical graphics and mapping package that combines two kinds of expertise data analysis and data presentation. The program has a repertoire of over twenty chart types with 3-5 data variations for each. The data analysis portion of the program has knowledge about the types of charts that are appropriate for different

kinds of statistical analyses. The data presentation portion of the program has knowledge about standard charting practices in different professional disciplines and knowledge about effective graphic design of charts which it embodies in default style specifications.

As the "formatting intelligence" of charting and diagramming packages grows, the need will arise for users to coordinate the style of graphics produced by their charting program with the graphical style of their page makeup software. In some applications this will never be an issue. In others, the need to coordinate may favor vendors whose packages play well together.

AI-based typography

AI-based techniques are being used to generate typefaces. For example:

Bitstream (Cambridge, MA) has developed an intelligent typeface scan conversion program that reads hand-drawn masters and creates optimal bit-map representations at any size and any resolution.

Compugraphic (Wilmington, MA) has an OEM product called Intellifonts that generates screen bit-map typefaces in actual size from outline masters on the fly.

The direction that we are heading appears to be toward "intelligent typefaces" that contain information about how to render themselves in different circumstances.

This same idea may extend to clip art (smart clip art) and other reference graphics that are sold commercially.

Publication knowledgebase management

Publication knowledgebase management is concerned with building, maintaining, and accessing document libraries and multimedia knowledge structures.

One of the keys to the knowledge-intensive systems of the 1990s will be tools to intelligently structure, interrelate, and manage large amounts of information.

Personal, group, corporate, and commercial knowledgebase management tools will be a major focus of development.

Hyperpublishing technologies are being used to:

Automatically index and cross-reference information

Construct dictionaries and thesauruses

Semantically network relevant topics and subtopics

Intelligently search and retrieve

Network cooperative hypermedia knowledgebases

Compress/decompress imagery, text, and graphics

Desktop knowledgebases

Over the next few years, we are likely to see a trend from files and databases to knowledgebases as the foundation for electronic publishing systems.

Knowledge workers collect and abstract relevant information from large volumes of highly unstructured and rapidly changing material (much of it textual). No existing spreadsheet or relational database is of much use in helping knowledge workers with these tasks.

Despite AI's focus on knowledge-based expert systems, no practical way exists currently to support the creative abstraction and evaluation of unstructured information and the decision making involved in communicating and applying this new knowledge.

According to Chris Locke (Intelligent Technology Group, Pittsburgh, PA), the direction we are heading is toward desktop tools that can parse and tag document content, automatically index topics, construct a thesaurus, and semantically network relevant topics and relationships. These systems will also intelligently classify text-based knowledge elements, automatically create hypermedia linkages, search and retrieve similar documents, extract relevant text, and prepare abstracts, precis, and summaries.

Automatic indexing

Document understanding involves reading documents by computer and assimilating the information that the documents contain into a much larger framework of knowledge.

Topic structure is very important to effective search and retrieval. A key problem to be solved as we absorb new sources of information from external database services is how to structure our knowledgebase so that we will be able to correlate it with outside sources.

Most database services characterize the content of bibliographic item by assigning subject headings or classification codes.

These systems require a controlled vocabulary of keywords organized in a thesaurus and a classification scheme such as the Dewey Decimal System or the Library of Congress Subject Headings.

Electronic librarians

Electronic librarians are expert systems for browsing, search, and retrieval that can provide unlimited access to libraries and databases.

People who plan database services say we are heading toward "electronic librarians" and that the logical place to start is by making intelligent front-ends to existing database services.

Electronic librarians will help select appropriate sources, formulate search strategies, suggest alternatives, gain access to local and remote databases, follow correct local request procedures, extract relevant information, and index results. They will have metaknowledge about networks, remote databases, and search techniques and procedures in order to be able to interrogate various information sources at minimum costs. Electronic librarians will also have understanding of natural language in order to prepare precis and abstracts, and a knowledge of subject classification methods in order to index and cross-reference materials.

Intelligent search and retrieval

Most information retrieval packages use keywords and boolean logic in searching. This is only one way to perform a search, and for most people, it is not the most effective way. Many people, for example, have difficulty structuring boolean searches.

Another way to search is by similarity. In this technique, the user submits a fragment of text or even an entire document to the system and says, in effect, "give me more stuff like this." The system compares the content of the user's query with the available documents searching for similarity of topic and intensity. It ranks those documents according to how closely they appear to be relevant to the original query.

Executive Technologies' (Birmingham, AL) SearchExpress is an advanced search and retrieval package for the PC platform. It can handle text, spreadsheets, images, sounds, and even executable programs. The program creates an index for the "documents" or objects that it processes. Boolean operators are used to define search fields. The user can also reference the results of previous searches. The "search for similar documents" feature uses a weighted 20 word search, ignoring noise words (a, the, and, etc.) and words that occur frequently. The resulting list ranks objects by similarity. (It also takes time.) The "object linking" feature implements two-way hypertext-like links between files. Linkage, however, cannot be established between components of these files (words, graphics).

Third Eye Software (Palo Alto, CA) offers an intelligent search and retrieval product called Elixir that implements similarity searching in a file server.

Visual knowledgebases

Little work has been done to develop techniques for visual information management.

Development is needed. Most practical applications relate to textual information and bibliographically structured databases. Visual classification schemes, shape grammars, indexing, and intelligent search and retrieval techniques are still in a very early stage.

Hyperdocument generation

A great deal of attention has focused recently on the problem of intelligent document formatting and delivery of information through new information products.

Hyperpublishing technologies are being used to:

- Develop smart templates
- Intelligently format documents
- Intelligently retrieve text
- Format hypermedia publications
- Deliver expert communications
- Deliver text via speech synthesis

Smart templates

Vendors are starting to think about packaging graphic design expertise into smart templates so that authors will be able to develop acceptably formatted publications without a great deal of effort.

A smart template (or intelligent document format) contains graphic design elements and style specifications plus knowledge about how the design should be applied in various circumstances. Smart templates will be used with desktop systems combining wordprocessing, graphics, and AI-based tools (expert system) for document composition.

Currently, several companies are believed to be working on expert systems and "smart template" approaches for desktop publishing. For example,

Aldus (Seattle, WA) is rumored to be significantly enhancing the template features of PageMaker.

Archetype (Boston, MA) has been working toward an AI-based graphic design environment that would learn

designs by example. According to Archetype's president, Paul Trevithick, a key problem is to find a way for a graphic designer to accomplish AI-programming without having to program.

Cognition Technologies (Cambridge, MA) is developing an expert system for page layout that runs on the Macintosh called PageSmarts. The product can create layouts for reports or newsletters by guiding the user through the steps for producing the layout. From the information that it receives, it creates a template which can be used again.

Intelligent document formatting

Intelligent document formatting systems are starting to appear that compose documents directly from a database using a document profile and style specification. The new approaches use expert systems techniques and build on earlier "batch pagination" programs.

Expert Technologies Inc. (Pittsburgh, PA) has been developing expert systems for electronic publishing applications. One application is an expert system for telephone book layout and pagination that ETI has been prototyping for Nynex.

Crosfield/CSI (Glenrock, NJ) developed a sophisticated expert system for newspaper pagination. This project they have since discontinued.

Intelligent text retrieval

Industry and government users are moving toward intelligent user-friendly front-ends to help people access, interact, and create documents from large databases.

For example, when a user needs to call upon a database, the database may be linked to an expert system for intelligent data manipulation, and to another expert system (on the mainframe or on the desktop) to format a report.

Intelligent text retrieval projects have been undertaken by organizations such as AT&T, Citicorp, McGraw Hill, US West, Unisys, DEC, The Defense Intelligence agency, and many others.

Dow Jones News/Retrieval is implementing its stock quote and text search services using a Connection Machine to greatly speed search and retrieval times.

The Wall Street Journal has implemented its chart database in HyperCard and claims that "even the most unknowing freelancer can navigate the chart archive." In conjunction with the visual database they automated chart production through Microsoft Excel and Chart, and Adobe's Illustrator. Now, instead of transmitting graphics by facsimile, they transmit using PostScript.

Hypermedia publications

Corporations and government agencies are beginning to create hypermedia publications for training, maintenance, and general information.

The Air Force (Wright-Patterson AFB), for example, is developing hypertext aircraft maintenance manuals with variable display/print formats. The system is also capable of configuration control and version control so that original versions of the documentation may be kept intact while creating newer versions of the material.

Universities are using hypermedia to deliver tutorials.

For example:

Cornell University (Ithaca, NY) has implemented a hypermedia tutorial that teaches pathology called HyperPath.

Brown University (Providence, RI) has developed a complete hypermedia tutorial on Plant Cell Biology using its Intermedia system. The Intermedia system is a framework for a collection of tools that allow authors to create links to documents of various media such as text, timelines, diagrams, computer-generated images, video documentaries, and music.

University of Southern California (Los Angeles, CA) has developed a HyperCard-based hypermedia publication called the Anthropologist's Notebook. Hypercard is used to access a dynamic simulation of the ecological effects of various water management decisions in a southeast Asian country. Graphical displays are generated on-demand to depict the results of changing variables in the simulation. Menus and explanatory text can be displayed in several languages.

Expert communication systems

Expert communication systems are emerging as a new form of publishing that combines computer advisory dialogues, hypermedia, and subject-domain expert systems.

The intelligent interface makes it easier to control a system, access information the user wants, and present it in an appropriate form.

DARPA (Arlington, VA) is converting textual regulations into a knowledge base, building an expert system for accessing and delivering knowledge about what you can and cannot do according to the regulations. The system delivers the output in the form of hypertext.

Hughes Aircraft is combining AI and hypertext to help technicians troubleshoot complex weapon systems problems in the least amount of time. The troubleshooting assistant uses model-based (simulation) and rule-based approaches to diagnose. Hypertext is used to deliver (or browse) technical information.

Knowledge Garden's (Nassau, NY) KnowledgePro is a PC-based tool for developing expert communication systems. Communications with expert systems are often overly controlled while hypertext environments can be too flexible for handling directed inquiries. KnowledgePro offers a hybrid strategy. It combines hypertext (down to the word level), an expert system shell, and a procedural language into a development environment for communicating expertise. It allows the designer of the expert communications to ask questions of the user about his hypertext selections and then control the level and direction of information presented in response.

Speech synthesis

Speech synthesis is becoming a viable way to deliver information.

According to a study by Link Resources Inc. (New York, NY), sales of interactive voice services and equipment will grow at a compound annual rate of 32 percent through 1991 (from \$1.5 billion 1987 to over \$4.1 billion). Applications to which they refer include voice messaging, interactive voice response, audiotex, autodialers, integrated voice/data terminals, talking yellow pages, subscription services such as stock quotes, intelligent 800 services, 976 dial-it services, and 900 polling services. The key trend in equipment, they point out, is that single

functions, for example, just voice mail or just audiotex, is a thing of the past. Trend is toward multifunction systems.

Speech synthesis converts text into spoken language.

A significant number of companies are developing speech synthesis products: for example,

First Byte (Long Beach, CA) sells Smooth Talker, a text-to-speech synthesizer software program for the Apple Macintosh.

Votrax International (Troy, MI) manufactures a phoneme-based, text-to-speech system built around proprietary chips that contain 64 phonemes and provide control over speech rate, pitch, pitch movement rate, and amplitude. The speech synthesis system is also capable of comprehensive sound effects and music, generated independently of the speech so that speech, music, and sound can be combined.

Speech Plus (Mountain View, CA) offers a text-to-speech synthesis product that uses a MIT-developed simulation of human vocal tract and linguistic rules to produce extremely high quality speech.

Vynet (Los Gatos, CA) markets a voice-response system (8-24 lines) and development facility to digitally record, edit and update phrases, words, and voice prompts to the caller.

Granite Systems (Lafayette, CA) offers a voice-response system that takes messages via touch-tone input and plays pre-recorded human speech; constructs questions and answers from a lexicon of sentence fragments and words; and can initiate outgoing calls in response to incoming calls.

Publishing/printing management

Publishing/printing management is concerned with planning, scheduling, coordinating, and controlling publishing and printing operations in the graphic arts.

As an information product manufacturing process, publishing and printing will afford numerous opportunities to apply expert systems and other AI-based approaches to the problems of managing production, printing, and distribution of published information.

Hyperpublishing technologies are being used to:

Manage the development of documents and publications

Provide quality color separations

Configure presses

Monitor print quality on press

Rockwell International (Chicago, IL), for example, has developed an expert system for the Los Angeles Times that determines a suitable location on press for each page image plate of the newspaper. The program contains specific knowledge about the press, the configuration procedure, and the physical and operational constraints surrounding the press.

Next Issue

Color system vendors

The emerging color systems market has new rules and new players. Our next issue brings you a comprehensive survey of the color system vendors.

DAVIS REVIEW

C A L E N D A R

June 1988

June 5-8
IABC's Annual Conference
Anaheim, CA
International Association of Business
Communicators
415/433-3400

June 6-8
Expert Communication '88
San Jose, CA
Graphic Communications Association
703/841-8160
First conference devoted to AI in publishing

June 6-8
Computer Graphics for Design/Spring'88
New York, NY
Pratt Center for Computer Graphics in
Design
914/741-2850
Major conference on electronic design

June 11-15
ANPA/TEC '88
Atlanta, GA
American Newspaper Publishers
Association
703/648-1211
Newspaper technology and systems

June 12-17
**International Design Conference:
The Cutting Edge**
Aspen, CO
IDCA, P.O. Box 664, Aspen, CO 81612
Annual design conference

June 13-15
Magazine Publishers Congress
New York, NY
Magazine Publishers Association
212/752-0055

June 13-15
Gate 2000: Reaching Mass Markets
Washington, D.C.
Videotex Industry Association
703/522-0883

June 14-17
Neocon 20
Chicago, IL
The Merchandise Mart, Chicago, IL 60601
World congress on environmental planning
and design

June 18-21
**National Convention of the American
Advertising Federation**
Los Angeles, CA
202/898-0089

June 21-23
PC Expo
New York, NY
201/569-8542
PC show for resellers & volume buyers

June 23-25
Type-X '88
Orlando, FL
Typeworld
603/898-2822
Composition and electronic publishing
technology

June 26-28
**Corporate Electronic Imaging Markets
Conference**
Bedford, MA
Institute for Graphic Communication
617/267-9425

July 1988

July 2-8
Ausgraph '88
Hyatt on Collins Melbourne, Australia
Ausgraph 88 Conference Secretary
P.O. Box 29 Parkville Victoria 3052
Australia

July 4-8
State of the Art in Computer Graphics
Dansford, England
Association for Computing Machinery
44/532-431751

July 8-13
PPA '88
Orlando, FL
Professional Photographers of America
312/299-8161
Official convention and trade show for
professional photographers

July 18-22
Video Expo Chicago
Expo Center, Chicago, IL
Knowledge Industry Publications
914/328-9157
Video technology

July 21-23
Desktop II
Nashville, TN
National Newspaper Association
202/466-7200

July 24-27
IEEE Neural Networks
San Diego, CA
IEEE
619/222-7447
Emerging AI technology

August 1988

August 1-5
SIGGRAPH'88
Atlanta, GA
ACM/Siggraph
312/644-6610
Premier computer graphics conference

August 7-10
IPHC'88 Annual Conference
Cincinnati, OH
International Printing House Craftsmen
513/891-0611

August 11-13
MacWorld Exposition
Boston, MA
Bayside Expo Center
617/329-7466
Major exposition of Macintosh technology

August 15-20
AMI'88 Conference
Phoenix, AZ
Association for Multi-Image
813/932-1692

August 21-26
AAAI '88 National Conference
St. Paul, MN
AAAI
415/328-3123
National conference on artificial
intelligence

August 24-27
WORLDDESIGN'88/NEW YORK
New York, NY
Industrial Designers Society of America
703/759-0100
Annual industrial design conference

DAVIS REVIEW

*Electronic Design and
Communication Arts*

Subscription Rates: \$300 in US, \$325 in Canada and Mexico,
\$350 other than N America. (US currency only)

☐ Enclosed is my check in the amount of \$_____ for a full
year's subscription of DAVIS REVIEW. (Ten issues per year mailed
first class to US and Canada, airmail overseas)

Make checks payable to DAVIS REVIEW. You may transfer funds
directly to our bank: National Bank of Washington, Washington, DC
20036 into account number: 1 009 62 1. If you transfer funds bank
to bank, be sure to identify the name of the subscriber.

Name _____

Company _____

Title _____

Address _____

City, State, Zip _____

Country _____

Telephone _____

Mail to: DAVIS REVIEW 2704 Ontario Rd NW Washington, DC 20009
Telephone: 202-667-6400 **Telex:** 5106018770 DAVIS